Protecting Your Assets: Information Security in the Teradata Database

03.15 EB7075 SECURITY

TERADATA.

Information Security Awareness

Companies are driven by information today. Leading companies rely on information to make routine operational decisions throughout the day in addition to major strategic decisions which will guide the organization's future. It's no wonder that information has become one of the most valuable assets within most companies.

Unfortunately, criminals have also realized that data is valuable. And, because it can be stolen from a distance, without personally walking into a branch or retail location, cyber-crime is increasing substantially. A Ponemon Institute study identified six data breaches in 2014 that exposed data about 1.1 million up to 145 million people each. Another Ponemon study found that the average cost of cyber-crime continues to rise for companies, and the time to resolve a cyber attack is also rising. Consider this: this study was Ponemon Institute's fifth annual Cost of Cyber Crime Study. This illustrates how commonplace it has become and the engrained impact of cyber-crime and information security on today's companies.

Effective information security builds and protects corporate brands and maintains customer loyalty. Consumers have come to understand that companies collect information for better customer service and marketing. Savvy consumers understand that they "exchange" personal data for better products, services, and prices, but only do so with companies they trust to be good stewards of their personal data.

Defense-In-Depth

A single product, feature, security mechanism, or process cannot be relied upon in today's business world. Building a strategy with many dimensions provides a higher level of security and risk mitigation. Defense-in-depth is a strategic roadmap for implementing multiple layers of security defense mechanisms to ensure protection of business critical information in the event a single mechanism fails or is compromised. The dimensions are far ranging from technological controls implemented to secure the database, operating system, server, and network access to the security procedures around managing vulnerabilities, and the auditing and monitoring of the system. Even using security policies and processes to certify and maintain the secure configuration are included.

Teradata advocates a defense-in-depth approach to information security. Combining this approach with an integrated data warehouse (IDW), which in itself reduces the dispersion of data storage throughout an organization, increases the overall effectiveness and manageability of a comprehensive security program. When data is in an IDW, security efforts can be focused and easily monitored and maintained without the risk of weaknesses in a single uncontrolled database exposing the entire enterprise environment to risk.

Available Product Capability and Services

Teradata has long recognized the need for comprehensive information security and designed products to enable customers to implement and manage a secure environment. These products have a variety of features and capabilities to implement security policies. Teradata also partners with an ecosystem of companies with a wide range of expertise and products to meet various requirements. Beyond products, best practice and information resources are available for guidance and Professional Services by Teradata security experts are offered to assess, plan, and implement security policies and measures.

Examples

The following articles illustrate some of the security tools that are available for use with the Teradata Database.

Secure Access

Teradata Row Level Security better safeguards who sees your data.

See page 4

Tokenization on the Node

Non-formula data security solution provides greater protection against cyber-criminals.

See page 6

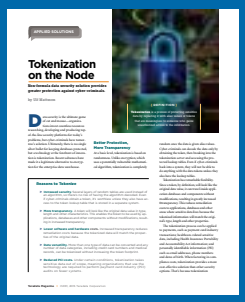

Security Features in Teradata Database

The Teradata Database supports many important features that are designed to enhance the security of an integrated data warehouse.

See page 8

10000 Innovation Drive, Dayton, OH 45342 **Teradata.com**

03.15 EB7075

TERADATA

Secure Access

Teradata Row Level Security better safeguards who sees your data.

by Jim Browning

Traditionally, implementing row-level security in the Teradata Database has been by using views, where a ccess rules are encoded using SQL. These rules can simply filter data based on a predicate evaluation. They can also be more complex, joining one or more security tables that may supply attributes of the users or applications for an evaluation.

Teradata Row Level Security, an alternative method, is designed to meet the requirements for Mandatory Access Control (MAC) while allowing for extensibility to meet other row-level security nee ds, particularly those similar to MAC. However, it is not a replacement for view-based approaches, since those may often be better suited to the many applications that require the ability to encode filtering rules in SQL and use information from different tables.

Better Filter System

Teradata Row Level Security, introduced in Teradata Database 14.0, provides filtering of table rows using attributes assigned to the user and related attributes defined for the table. It overcomes two specific problems common to view-based row-level security:

1. With the view-based approach, only the underlying data for access is secured through the views. Any other access to the underlying base tables effectively bypasses this security policy. But Teradata Row Level Security is applied to the base tables and the security policy is enforced on any access to those tables.
2. Filtering access for SELECT operations works very well with the view-based approach, but views usually cannot be used to filter access for INSERT, UPDATE and DELETE operations. Teradata Row Level Security provides the capability to filter access for all these.

Teradata Row Level Security is based on user-defined policy functions assigned to database tables, and it filters access to rows whenever data in the tables is queried or updated. A new type of database object, called a CONSTRAINT, identifies the policy functions and the attributes to be evaluated by those functions for row-level security filtering.

New access rights restrict the management of the row-level security policies to one or more designated database security administrators. With these rights, the administrator can:

- Create and manage CONSTRAINT objects
- Assign CONSTRAINT names to users and profiles
- Assign CONSTRAINT objects to tables

SECURITY FILTER

Row-level security refers to the *filtering* of access to rows within a database table. Use of row-level security enables data maintained in a common set of tables to be accessed by different users when each may be allowed to see and write to only a subset of the data based on security requirements or policies.

Grant and Deny Data Access

The Teradata Database implements a typical Discretionary Access Control (DAC) security model. With this model, the owner of a database object can grant access to users and roles as desired.

A second model, MAC, restricts access to database objects based on data sensitivity and the authorization level of the user—commonly required for securing highly sensitive data in government agencies around the world. The basic MAC model defines a relationship between objects (data) and subjects (users and applications). Data within an object (i.e., rows within a table) can be assigned a classification to indicate its sensitivity. For example, a government organization could classify data using hierarchical levels such as Top Secret, Secret, Confidential and Unclassified. Subjects are subsequently assigned a similar clearance that determines their level of access.

Classifications and clearances are often described as security labels, which usually include two components:

- A required hierarchical component that is used to indicate the sensitivity of the data
- An optional non-hierarchical set of compartments that can be used to identify and restrict access by category or group

With Teradata Row Level Security, CONSTRAINT objects are used to construct these labels, which are then assigned to objects and subjects. They are subsequently used to enforce access rules.

MAC enforces rules to restrict access to labeled data. The rules define who can read and write the data. The enforcement of the MAC policy does not replace the standard DAC policy. The Teradata Database first enforces DAC, ensuring that the user has appropriate access rights on the table, before enforcing MAC.

Better Approach to Sensitive Data

Teradata Row Level Security is primarily designed to meet requirements for MAC, which determines access to data based on its sensitivity and the user's authorization level. Teradata Row Level Security enhances table- and column-level security capabilities by giving users access only to relevant data, even if that means they can only see certain rows in a table. The level for users can be as deep or layered as necessary. These security measures make sure users do not see or write to data in protected tables without proper clearance. T

Jim Browning is the enterprise security architect for Teradata Labs.

ONLINE

For technical details on using row-level security and examples that can be used as a template, see the Teradata Row Level Security Orange Book on **Teradata.com.**

APPLIED SOLUTIONS

Tokenization on the Node

Non-formula data security solution provides greater protection against cyber-criminals.

by Ulf Mattsson

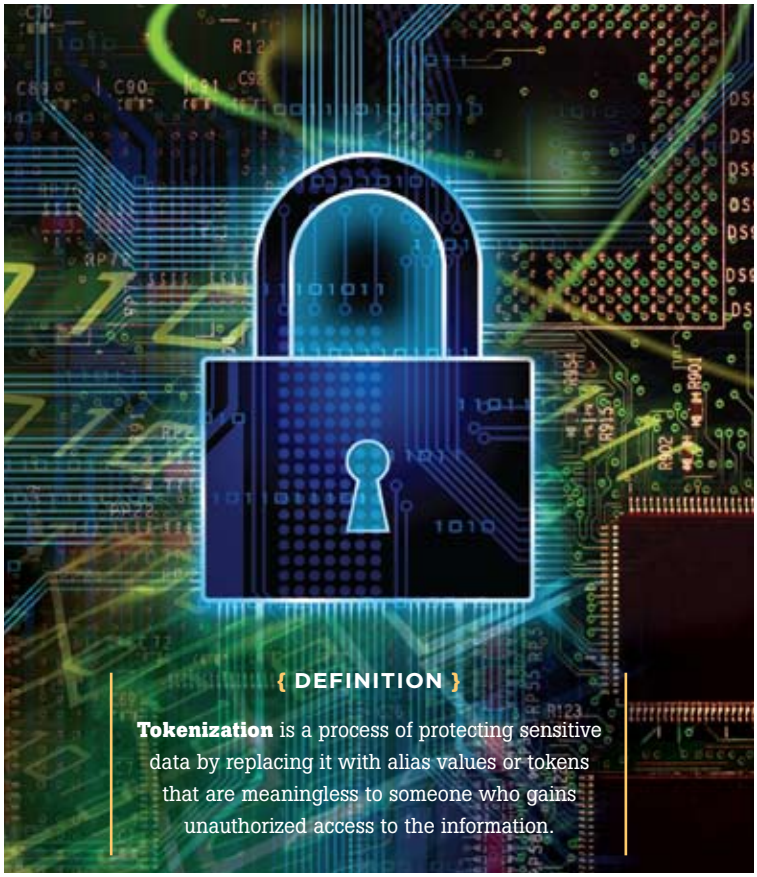

{ DEFINITION }

Tokenization is a process of protecting sensitive data by replacing it with alias values or tokens that are meaningless to someone who gains unauthorized access to the information.

Data security is the ultimate game of cat and mouse—organizations invest countless resources researching, developing and producing top-of-the-line security platforms for today's problems, but cyber-criminals have tomorrow's solution. Ultimately, there is no single silver bullet for keeping databases protected, but a technology at the forefront of innovation is tokenization. Recent advances have made it a legitimate alternative to encryption for the enterprise data warehouse.

Better Protection, More Transparency

At a basic level, tokenization is based on randomness. Unlike encryption, which uses a potentially vulnerable mathematical algorithm, tokenization is completely random once the data is given alias values. Cyber-criminals can decode the data only by obtaining the token, then breaking into the tokenization server and accessing the protected lookup tables. Even if cyber-criminals hack into a system, they will not be able to do anything with the data tokens unless they also have the lookup tables.

Tokenization has remarkable flexibility. Since a token, by definition, will look like the original data value, it can travel inside application databases and components without modifications, resulting in greatly increased 0transparency. This reduces remediation costs to applications, databases and other areas where sensitive data lives because the tokenized information will match the original's type, length and other properties.

The tokenization process can be applied to payments, such as payment card industry transactions; healthcare-related sensitive data, including Health Insurance Portability and Accountability Act information; and personally identifiable information (PII) such as email addresses, phone numbers and dates of birth. When factoring in compliance costs, tokenization provides a more cost-effective solution than other security options. That's because tokenization

Reasons to Tokenize

- **Increased security.** Several layers of random tables are used instead of an algorithm, so there's no risk of having the algorithm decoded. Even if cyber-criminals obtain a token, it's worthless unless they also have access to the token lookup table that is stored in a separate system.
- **More transparency.** A token will look like the original data value in type, length and other characteristics. This enables the token to be used by applications, databases and other components without modifications, resulting in increased transparency.
- **Lower software and hardware costs.** Increased transparency reduces remediation costs because the tokenized data will match the properties of the original data.
- **Data versatility.** More than one type of data can be converted and any number of data categories, including credit card numbers and medical records, can be tokenized without increasing the token footprint.
- **Reduced PCI costs.** Under certain conditions, tokenization takes sensitive data out of scope, meaning organizations that use the technology are required to perform payment card industry (PCI) audits on fewer systems.

provides better separation between the data and the security system. This can take the sensitive data out of scope, eliminating the requirement and costs to meet compliance. For organizations that need to have the actual data stored within their company, tokenization greatly reduces the extent and use of sensitive information, dramatically reducing compliance costs.

Next-Generation 'Small Footprint' Tokenization

First-generation tokenization solutions are characterized by their use of large token servers that constantly grow with encrypted data. The next generation of tokenization is defined by a small system footprint that compresses and normalizes the randomized data. This implementation can be fully distributed and does not require any resource-intensive data replication or synchronization. Users benefit from a lower-cost, less-complex, higher-performing data security solution that guarantees no collisions.

Small-footprint tokenization offers advantages over first-generation methods:

- **Latency.** Token servers with small footprints enable the distribution of the tokenization process so that operations can be executed in parallel and closer to the data. Therefore, latency is eliminated or greatly reduced, depending on the deployment approach used.
- **Replication.** The smaller footprint allows the creation of farms of token servers that are based on inexpensive commodity hardware. This permits any required scaling without the need for complex or expensive replication processes, which eliminates collisions.
- **Versatility.** Any number of data categories, including credit card numbers and medical records, can be tokenized without the challenges that previous tokenization methods impose. More data types can also benefit from the transparent properties that tokens offer.

Six Tokenization Myths

There are many misconceptions about tokenization. Here are six pervasive myths and the facts to dispel them:

MYTH 1: **IT IS A NEW DATA SECURITY SOLUTION.**
Tokenization was first introduced in 1300 by the Vatican and reintroduced in 1998 by IBM, so it's a technology that has been around for a long time. However, it has recently become a subject of widespread interest by the financial service industry, and the Payment Card Industry Standards Security Council is expected to release guidelines on tokenization this year.

MYTH 2: **IT IS ONLY FOR CREDIT CARD DATA.**
Tokenization is a viable solution for multiple types of sensitive data, including credit card numbers, social security numbers, email addresses, etc. Tokenization limits the exposure of sensitive information, minimizing the chance that a hacker will obtain any useful data.

MYTH 3: **ENCRYPTION ISN'T USED.**
Sensitive data is encrypted in Compliance 3 when "resting" as a preventive measure in case the data center is hacked. Encryption is used on the token server to protect the sensitive data.

MYTH 4: **THE NEED FOR PCI COMPLIANCE IS ELIMINATED.**
For *most* organizations, tokenization can reduce payment card industry (PCI) scope and/or the need for compliance. For example, a small organization may never need to store sensitive data on its servers because a token can provide sufficient payment information to complete a transaction without the need for the original data. For large retailers and financial institutions, tokenization would narrow the PCI scope, but they would still need to handle the original data, which would necessitate compliance on some systems.

MYTH 5: **TOKENIZATION IS ALWAYS THE BEST SOLUTION.**
The innovative technology will change how people think about safeguarding sensitive data. It's an excellent solution for protecting credit card numbers, Social Security numbers, health records and personal information such as email addresses, passwords and birth dates. However, encryption is a better option when protecting files such as scanned documents and image and sound files.

MYTH 6: **IT CREATES COLLISIONS.**
Token collisions, also known as replications, have been a major concern. Current tokenization solutions have been tested and proved to fully eliminate the risk of collisions.

Tokenization can apply to all types of enterprises and is an especially good fit for the Teradata Database. When deployed on the node, small-footprint tokenization can be fully distributed across nodes and AMPs on a Teradata system.

Data Protection for Today and Tomorrow

With its transparency characteristics and ability to protect sensitive information, tokenization will become a more popular security solution because of standards by the Payment Card Industry Standard Security Council and data protection requirements that go beyond payment data. Standards will ultimately help set best practices and govern tokenization implementation. While tokenization doesn't solve every data security conundrum, it will protect data today and for many years to come, since there is no encryption algorithm or key that must be refreshed periodically to prevent the growing threats of malicious attacks. T

Ulf Mattsson is the chief technology officer for Protegrity and created the architecture of the Protegrity Data Security Platform, which includes tokenization.

Security Features in the Teradata Database

Jim Browning

02.15 EB1895 SECURITY

TERADATA.

Table of Contents

Abstract

The Teradata® Database supports many important features that are designed to enhance the security of an integrated data warehouse. These security features include:

- User-level security controls
- Increased user authentication options
- Support for security roles
- Enterprise directory integration
- Network traffic encryption
- Auditing and monitoring controls

This white paper provides an overview of the security features and describes scenarios for their usage. The paper will also discuss the independent evaluation of the Teradata Database to the international Common Criteria for Information Technology Security Evaluation (ISO 15408) standard.

Security Challenges

Increased public attention to security is driving the restructuring of security requirements. The role that IT will play in helping address these challenges will be significant. However, IT departments are under pressure to cut their operating costs, while being asked to improve and standardize information security. Teradata's security approach will assist Teradata Database Security Administrators who are facing these new challenges.

Legislated requirements, government regulations, and industry standards all result in a continually evolving security landscape. The following are examples that are driving increased requirements for data warehouse security across many industries and geographies:

European Union Privacy Directives

The principles established by the European Union (EU) Privacy Directives serve as the foundation for many international privacy and security laws. These directives require the use of appropriate technical and organizational measures to ensure confidentiality and security of processing of personal data.

Health Insurance Portability and Accountability Act

The Health Insurance Portability and Accountability Act of 1996 (HIPAA) mandates standards and requirements for maintenance and transmission of health information that identifies individual patients and compliance is required by all U.S. health care organizations that maintain or transmit electronic health information. A Security Rule establishes specific security requirements for authorization, authentication, audit trail requirements, secure data storage and transmission, and data integrity.

Gramm-Leach-Bliley Act

The Gramm-Leach-Bliley Act of 1999 (also known as the Financial Modernization Act) requires that financial institutions adopt of policies and procedures to provide for the protection of financial information that identifies individual consumers. Such procedures must protect against any anticipated threats or hazards and protect against unauthorized access which could result in substantial harm or inconvenience to a customer.

Sarbanes-Oxley Act

The Sarbanes-Oxley Act of 2003 includes a number of reforms intended to increase corporate responsibility, improve financial disclosures, and protect against corporate and accounting fraud. While this legislation does not mandate the use of specific security controls, Section 302 does require that internal controls be established to protect data from both internal and external threats, and Section 404 requires that corporations report on the effectiveness of those controls. Also, Section 409 requires the disclosure of any material changes to the financial condition or operation of the company (potentially to include a major security compromise).

Personal Information Protection Act (Japan)

The Japanese Personal Information Protection Law requires that companies operating in Japan develop and implement information privacy and security controls for any databases containing consumer or employee information on more than 5000 persons. Japan's Ministry of Economy Trade and Industry (METI) has issued specific guidelines for maintaining the security of these databases.

Payment Card Industry Data Security Standard

Developed by Visa and MasterCard, the Payment Card Industry Data Security Standard applies to merchants and service providers that store, transmit, or process credit card transactions. The standard outlines twelve specific requirements that must be implemented to protect cardholder information.

Security, as an aspect of IT control requirements, defines an attribute of information systems, and includes specific policy-based mechanisms and assurances for protecting the *confidentiality* and *integrity* of information, the *availability* of critical services and, indirectly, privacy. Data in a data warehouse must be protected at both ends of a transaction (user and enterprise). Figure 1 depicts the relationships in simple terms.

These concepts and relationships are taken from the Common Criteria ISO 15408[1] standard specifying the *"Privacy Class of Common Criteria."* It proposes that all security specifications and requirements should come from a general security context. This context states that "security is concerned with the protection of assets from threats, where threats are categorized as the potential for abuse of protected assets."

1 Common Criteria for Information Technology Security Evaluation, Part 1: Introduction and general model

Data warehouse security requires protection of the database, the server on which it resides, and appropriate network access controls. Teradata highly recommends that customers implement appropriate network perimeter security controls (e.g., firewalls, gateways, etc.) to protect network access to a data warehouse.

The remainder of this paper will specifically discuss some of the security features that can be employed to effectively secure a Teradata Database.

Teradata Database Security Features

Teradata is continuously adding security features to its products. We are committed to driving significant benefit for our customers now and into the future, and achieve our vision for a leadership role in data warehouse security.

The following sections describe some of the security features that aid Teradata Database clients in effectively implementing a data warehouse security policy, and highlight some attributes and intended usage of these features.

Authentication

Authentication refers to the process of establishing the legitimacy of a user before allowing access to database resources. Proper authentication of users is fundamental to ensuring the security of any database system. The Teradata Database provides multiple options for authenticating database users. Additionally, custom authentication methods can be developed and deployed to further enable integration of a Teradata data warehouse into diverse security management environments.

All supported authentication methods are described by a set of properties that can be managed by a security administrator. These properties allow for the security administrator to establish default authentication methods and to restrict or limit the methods that may be selectable by a database user. Other properties may similarly be managed by the security administrator.

User-level Security Controls

Typically, a database user must provide a valid username and password as part of the logon string in order for a database session to be established. However, properly securing such password-based schemes requires that a security administrator be able to ensure that passwords are regularly changed, are sufficiently complex, and that effective precautions can be taken to protect against attempts to guess user passwords. As such, the Teradata Database supports a rich set of password security controls that can be specified at either the user level or the system level. This is important since it is often desirable to establish and enforce different password management policies for different types of database users (e.g., batch versus interactive).

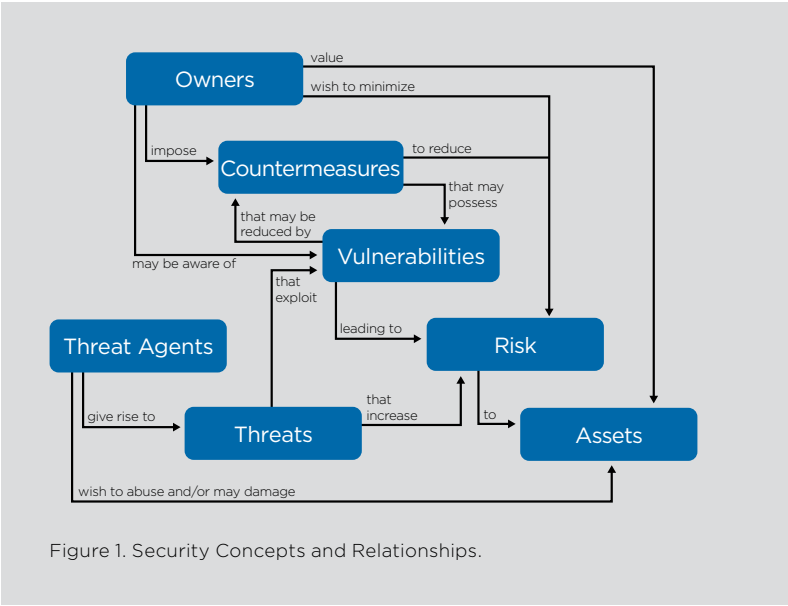

Figure 1. Security Concepts and Relationships.

User-level controls are implemented using the User Profiles feature. In this manner, profiles specifying specific password management policies can be defined and assigned to individual users, groups of users, or an entire enterprise. When a user logs on to the Teradata Database, any associated profile password controls will take effect. If no associated profile password controls have been defined, then the system-level controls will take effect.

Figure 2. describes the password security controls that are supported by the Teradata Database (reference the *Security Administration* reference manual for implementation specifics[2]).

2 Teradata Database Security Administration – www.info.teradata.com

Usage Controls	Description
Password Expiration	Allows the security administrator to define a time span during which a password is valid. After the time elapses, the user must change the password.
Password Reuse	Allows the security administrator to define the time span that must elapse before a previously used password can be reassigned to a user.
Maximum Logon Attempts	Allows the security administrator to define the number of erroneous sequential logon attempts a user is allowed before the user is locked to further logon attempts.
Password Lockout time	Allows the security administrator to sets the user lock time duration after the user has exceeded the maximum number of logon attempts.
Format Controls	**Description**
Password Length	Allows the security administrator to define the minimum and maximum number of characters required in a valid password string.
Password Construction	Allows the security administrator to specify whether alpha characters, digits, special characters, and a combination of upper and lower case characters are to be allowed or required in the password string. Also, allows the security administrator to specify whether the username or dictionary words should be allowed to be included in the password string.

Figure 2. Password Controls

Single Sign-on

Effective user authentication is a foundation of a database system's security services. However, secure authentication may be compromised in large, heterogeneous networks where users may be required to remember multiple user names and passwords. To address this issue, a single sign-on capability can be employed to allow network users to seamlessly access authorized network resources and applications, including an integrated data warehouse, with a single authentication that is performed upon initial network access. This capability improves the productivity of network users, reduces the cost of network operations, and, ultimately, improves network security. Further, security is improved by eliminating the need for an application to declare or store a password on the client system.

The Teradata Database supports single sign-on (SSO) of users that have successfully authenticated to a Windows domain or Linux Kerberos realm. With SSO, a user does not provide any credentials (user ID or password) when initiating a logon to the database. Rather, the user will be identified and authenticated based upon his/her current domain-authenticated session. Kerberos V5 is the protocol used to facilitate SSO of users to the Teradata Database. Figure 3 depicts the relationship between users, the Teradata Database server, and Microsoft Active Directory in implementing single sign-on.

LDAP Authentication

For enterprises where users may have access to many applications and systems, it is common to manage separate user accounts for each application resulting in

redundant and/or inconsistent data and increased user management costs. This lack of centralization also represents a significant security risk in that unused or expired accounts and privileges are subject to misuse. As such, many enterprises are adopting centralized security management frameworks that provide for a single point of administration for internal and external users, configuration information, and security policies. Such systems can often simplify the process of creating, modifying and deleting user accounts, as well as authorizing access to protected resources.

These systems typically store and manage user information through a directory service that supports the Lightweight Directory Access Protocol (LDAP). LDAP-enabled applications, services, and databases can readily leverage a single, centralized repository of user information to control user access.

The Teradata Database supports an LDAP authentication method that allows for authentication of database users against a centralized LDAP directory rather than using credentials maintained in the data dictionary. This method authenticates a user (by means of the user's distinguished name and password) through a secure LDAPv3 bind to the directory.

Trusted Sessions

The Teradata Database support a Trusted Sessions feature that can be used to preserve the security model commonly used by middle-tier and web-based application systems. This features provides the ability to authorize a limited trust model that allows middle-tier applications to assert user identities and roles through pooled connections. Eliminating the need to authenticate (or re-authenticate) users accessing the database through a middle-tier results in improved performance, scalability, and support for large numbers of users. It also enables greater granularity of access rights enforcement and auditing of access based upon end-user identity.

IP Filters

IP Filters can be defined to restrict the use of non-interactive logons (used for batch load jobs, middle-tier application connection pools, web services, etc.) by client system IP address or network subnet. Restricting the use of these user logons to specific client or server machines helps mitigate the associated security risk since a hacker would also have to compromise a specific machine in order to effectively logon to the database.

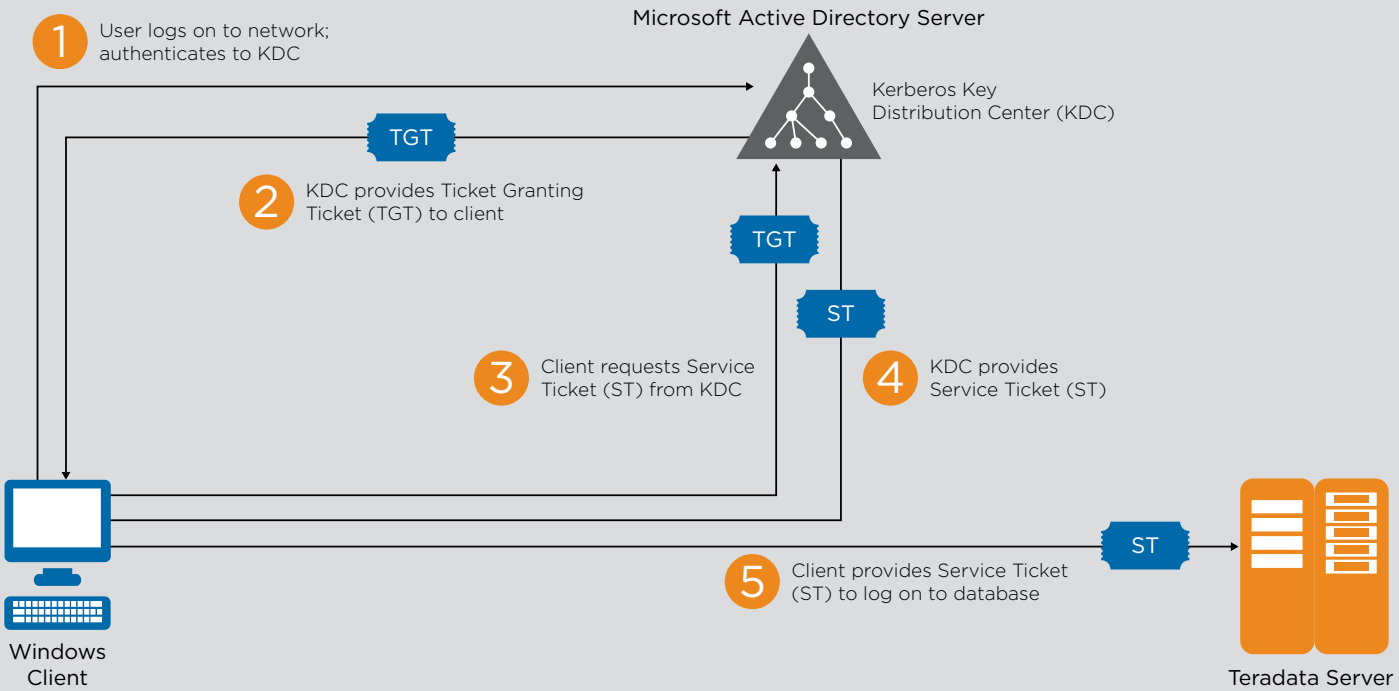

Figure 3. Kerberos-based Single Sign-on.

Authorization

Ensuring appropriate and authorized access to data is a major objective—and concern—in database security. The Teradata Database contains a robust set of fully integrated system access control capabilities. The mission of security administration on a Teradata Database system is to prevent unauthorized persons from accessing the system and its resources, as well as permitting legitimate users access to those resources to which they are authorized. The Teradata Database supports a discretionary access control policy in that access to database objects is restricted based upon the identity of users and/or groups to which they belong. The controls are discretionary in the sense that a user with certain access permissions is capable of passing those permissions on to other users.

Security Roles

One of the most challenging problems in managing large data warehouse systems is the complexity of security administration. Often, security administration is costly and prone to errors because security administrators must specify access controls individually for each database user. Role-based access control (RBAC) is a technology that can reduce the complexity and cost of security administration in large data warehouse environments. With RBAC, security is managed at a level that more closely corresponds to an organization's structure. Each database user may be assigned one or more roles with each role assigning access rights or privileges that are permitted to users in that role. Security administration with RBAC requires determining the operations that must be allowed by users in particular jobs and assigning those users to the proper roles. RBAC effectively manages complexities resulting from differing roles or hierarchies, thereby easing the task of security administration.

The Teradata Database provides support for Security Roles, which are used to define access privileges on database objects. For example, a user who is a member of a role can access the specific views for which the role has been granted appropriate access rights or privileges. For integrated data warehouses that provide access to many users, the use of roles will significantly simplify access rights administration and enhance overall security. A security administrator can create different roles for different job functions and responsibilities. For example, a security administrator can grant rights on a clinician view to a role and have these rights automatically applied to all users assigned to that role (Figure 4.).

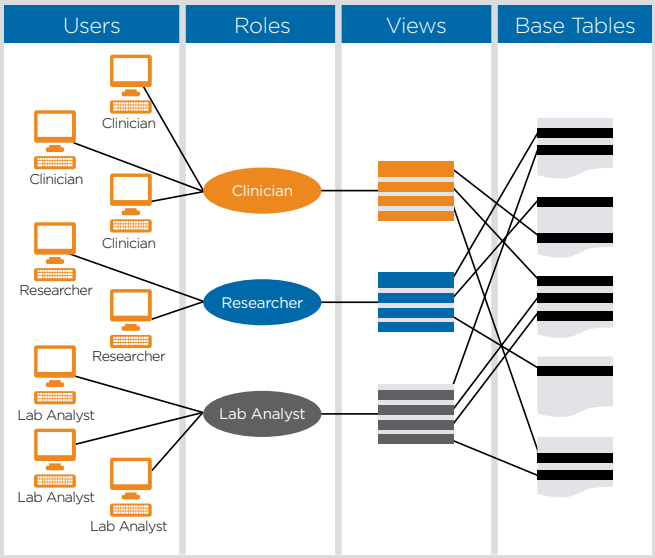

Figure 4. Security Roles.

Management of access rights is simplified by allowing grants and revokes of multiple access rights with one request. This is important when a user changes job functions (role) within the company. Should a job function need a new access right, it can be granted to the role and would be effective immediately for all users with that role.

To effectively use the Security Roles feature, individual rights must be converted into role rights. This requires creating the required roles and granting appropriate rights to each role. Roles can then be granted to users and users assigned their default roles. Finally, all individual access rights that have been replaced by role rights should be revoked from the users to ensure that all access rights are only granted through the role definition.

Typically, only one role will be the session's current or active role. Enabled roles are the current role plus any nested roles. At logon, the current role is the user's default role. Alternatively, it is possible to enable all roles granted to a user for a session.

Directory Integration

As noted earlier, many enterprises are adopting centralized security management frameworks, built using LDAP directory services, which provide for a single point of administration for users and associated security policies. Often, with such systems, the directory maintains access control policies that may be enforced by applications to authorize user access to enterprise resources.

Teradata has defined directory object mappings that allow for the mapping of the distinguished name of a directory user to a Teradata Database permanent user. Such users inherit the roles assigned to the mapped permanent user. However, additional external roles can be created and assigned to the directory user. External roles assigned to a directory user can be used in addition to any roles inherited from the mapped permanent user. A user profile may be created and assigned to a directory user in a similar manner. Upon successful authentication, the Teradata Database will enable the specified security role(s) and user profile for the database session.

The LDAP authentication method properties can be configured to allow for directory users that correspond to a user defined in the database to logon without requiring directory mappings. In this scenario, authorization to access database objects is managed entirely within the database.

Tools are provided to validate directory content and the operation of the directory when using the Teradata schema extensions.

Data Security

It is important to implement appropriate controls to protect sensitive data. Data can be vulnerable when transmitted over non-secure networks or when appropriate access controls have not been enabled for stored data. The Teradata Database provides facilities to manage the encryption of sensitive data when transmitted over non-secure networks. Further, row and column-level security can readily be implemented using database views.

Network Traffic Encryption

The Teradata Database and associated client applications and utilities typically operate in a traditional client/server environment. If clients are accessing the database server over non-secure networks, there is a risk that data may be compromised by a malicious user who is snooping on the network.

To mitigate this risk, the Teradata Database and client tools support encryption of data transmitted between client applications and the Teradata Database. Encryption is a CPU-intensive function that can negatively affect the performance of some operations. As such, its use should be carefully considered. The use of encryption is determined by the user through the client application and can be controlled on a per request basis. As such, the user has complete flexibility in the use of encryption to protect payloads transmitted over a network and to minimize any negative performance impacts. Alternatively, the client interfaces can be configured such that all sessions between the client applications and the database server are encrypted.

The security provided by encryption is dependent upon the strength of the encryption algorithm and the security of the key used to perform the encryption. The Teradata Database uses the public-key based Diffie-Hellman key agreement protocol to generate a secure key for use by the client and the database. A unique key is generated for each database session. The key generation is built into the underlying client-server communication protocol thereby eliminating the need for complex key management processes. Strong encryption is accomplished using the industry-standard Advanced Encryption Standard (AES) algorithm with support for 128-bit, 192-bit, or 256-bit keys.

In networked environments, a password transmitted from a client application to a database server may pose a security risk. If the password is transmitted in clear text over a non-secure network, there is a risk it could be intercepted by a malicious user snooping for data on the network. To protect against this, the Teradata Database client tools and utilities always encrypt the logon string (including username and password) that is transmitted to the Teradata Database server.

For compatibility purposes, the client and server are not required to be at the same version level. However, only the security features common to each version level can be used. This can allow for security features to be utilized according to individual client needs.

Full Disk Encryption

Some Teradata Appliances and Active Enterprise Data Warehouse systems support an optional full disk encryption capability implemented using self-encrypting drive (SED) devices within the storage arrays. Access to the data stored in an SED device requires an authentication key which is stored externally to the SED device and used by the storage initiator to "unlock" the device and enable READ/WRITE operations. Until a successful authentication has occurred the SED device will not respond to READ/WRITE requests. Once authentication has succeeded then the encryption key stored internally to the SED device is accessed and used by the device for encryption and decryption of the data stored and retrieved from the devices. Encryption is performed using AES-256.

Data Encryption

To protect sensitive data from internal and external threats, the Protegrity Database Protector for Teradata provides support for selective, and highly secure, column-level data encryption. A security administrator can centrally define, distribute, and enforce encryption policies independently from database administration with a full range of secure audit reports that provide complete accountability of access to secure data.

Data is encrypted using any of several supported industry-standard cryptographic algorithms in combination with initialization vectors to prevent statistical attacks and checksums on encrypted data to detect malicious attempts to manipulate the data. This capability provides centralization of all key management tasks on a single platform and effectively automates key management and administration, including an automated and secure mechanism for key rotation, replication, and backup.

The Protegrity Database Protector also supports tokenization which can often be used as an alternative to encryption for protection of sensitive data. This approach is governed by the same security policy and access control rules as used for encryption. Where tokens can be used in place of original data, this approach can eliminate much of the performance overhead commonly associated with the use of encryption.

The design fully exploits the parallelism provided by the Teradata Database to maximize database performance when querying encrypted data.

Teradata BAR Encryption

As part of Teradata's Backup, Archive, and Restore (BAR) architecture, encryption of backups and archives is supported through hardware-based encryption for LTO5 and LTO6 tape drives on Quantum i500 and i6000 libraries. Encryption is managed by two Scalar Key Management Appliances and is supported for backup-to-tape or copy-from-disk/tape. Encryption is performed using AES with 256-bit keys and does not affect the performance of backup or restore operations.

Encryption of backups to disk is supported using the EMC Data Domain DD4200 Deduplication Storage System. This system provides inline data encryption with compression. The encryption is performed using AES with 128-bit or 256-bit keys and implemented using FIPS 140-2 validated RSA BSafe cryptographic libraries.

Row and Column-level Security

The Teradata Database supports the granting and revoking of SELECT, INSERT, and UPDATE rights at the column level.

Database views are used to restrict the rows and columns that users (or groups of users) can access. Views are part of the SQL standard and can be thought of as virtual tables that can be accessed as if they were physical tables to retrieve data from the database. Views can be defined to reference columns or rows from underlying views and/or

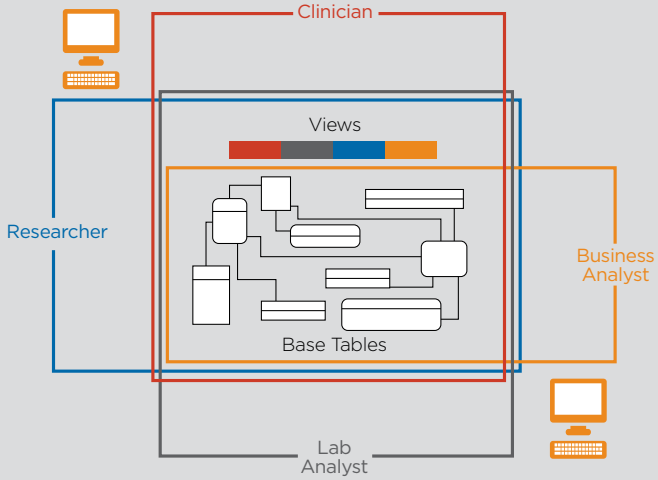

Figure 5. Database Views.

tables. A view does not actually contain data but rather is used to provide users with their own logical view of the data within the database. Figure 5 depicts an example from the healthcare industry where researchers, clinicians, lab analysts, and business analysts each represent a specific group of users with their own view of the database. These views enforce different security policies and access rights and privileges by limiting the data elements that are visible by each view.

Teradata Database support for views is particularly high performance in that the optimizer generates optimized SQL for selecting the appropriate columns and rows from the underlying base tables. Additionally, query access through views can generate very complex SQL expressions, which further exploit the inherent parallelism of the Teradata Database architecture.

The Teradata Row Level Security option can be used to restrict the rows of data that users (or groups of users) can access. It uses customer written policy functions to filter access based on the data in the row and the user submitting the query. Restrictions based on data sensitivity or ownership and access authorization can be easily enforced using Teradata Row Level Security.

Auditing and Monitoring

An important aspect of any security implementation is the creation and monitoring of a record of system activity to detect abnormal activity and to ensure that users are held accountable for their actions. To detect intruders and ensure data integrity, the Teradata Database provides a comprehensive set of auditing capabilities. A security administrator can periodically audit events on the Teradata Database to effectively detect potential attempts to gain unauthorized access to database resources or attempts to alter the behavior of the auditing facilities.

The Teradata Database automatically audits all logon and logoff activity. However, the security administrator can additionally configure the system's Access Log to log any successful and/or unsuccessful attempt to access any or all database objects by any or all database users. Also, the Access Log has controls to filter the logging by frequency of access or type of access. Teradata Database security features include the option to log the SQL expression that was used to perform the access to a database object. As such, all accesses are effectively audited.

Parameterized macros or triggers may be used to further customize or refine the auditing. Triggers are particularly useful when creating audit logs based upon specific data or content-based rules.

All audit information is stored in protected database tables within the data dictionary and access to the information requires appropriate access rights and privileges. The audit records can be viewed through ad hoc queries or with any appropriate application or query tool.

Assurance

Assurance refers to a level of confidence that a product's security features have been evaluated against a well defined and widely accepted set of security requirements. Security evaluations are conducted by independent, licensed and accredited organizations most often to the requirements of a specific industry standard. A security evaluation provides assurance through an analysis of a system's security functions using functional and interface specifications, guidance documentation, and the high-level design of the system to understand the security behavior. Independent testing of the security functions supports the analysis, evidence of developer testing based on a functional specification, selective independent confirmation of the developer test results, and a search for obvious vulnerabilities. Assurance is also provided through a configuration list for the system and evidence of secure delivery procedures.

Security Evaluation under Common Criteria

The Teradata Database has been independently evaluated to the requirements of the Common Criteria for Information Technology Security Evaluation (Common Criteria) standard. The Common Criteria is a multi-part standard that aligns with the International Standard ISO/IEC 15408:2005, which is meant to be used as a basis for evaluating security properties of Information Technology (IT) products and systems. The Common Criteria are defined by seven governmental security organizations known as "the Common Criteria Project Sponsoring Organizations" represented by Canada, France, Germany, the Netherlands, United Kingdom, the U.S. National Institute of Standards and Technology, and the U.S. National Security Agency.

The evaluation, conducted by the Science Applications International Corporation (SAIC) Common Criteria Test Lab under the National Information Assurance Partnership Common Criteria Evaluation and Validation Scheme (CCEVS), determined that the evaluation assurance level (EAL) for the product is EAL 4 augmented with ALC_FLR.3. The Teradata Database has been evaluated against 25 separate security functional requirements that describe the security behavior of the system. These requirements spanned multiple functional classes including Identification and Authentication, User Data Protection, Access, Security Audit, Security Management, and others. While the evaluation considered the design of the system, it also considered processes used for testing and installation and included a vulnerability analysis. As such, this evaluation provides a high level of assurance in the security design and implementation of a Teradata Database system.

This evaluation is intended to satisfy the requirements of those customers (primarily government agencies) that are required to procure only IT systems for which the security robustness has been formally evaluated and validated.

Conclusion

The Teradata Database provides a rich set of security controls for managing, protecting, and auditing access to stored data. These capabilities include extensive password controls, support for multiple authentication methods, access controls, high-performance database views, network traffic encryption, access logging, and audit reporting.

New industry regulations, especially in the retail, financial services, and healthcare industries, present increased challenges for securing an enterprise's information assets. The security capabilities described in this paper support many existing features within the Teradata Integrated Data Warehouse (IDW) and can assist Teradata Database security administrators in meeting these new challenges.

10000 Innovation Drive, Dayton, OH 45342 **Teradata.com**

02.15 EB1895

TERADATA.